IDA

INSTITUTE FOR DEFENSE ANALYSES

Evaluation of EarthRadar UXO Testing at Fort A.P. Hill

M. T. Tuley
J. M. Ralston
F. S. Rotondo
A. M. Andrews
E. M. Rosen

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

1. REPORT DATE (DD-MM-YYYY) 01-09-2001	2. REPORT TYPE Final	3. DATES COVERED (FROM - TO) xx-06-2000 to xx-07-2001

4. TITLE AND SUBTITLE
Evaluation of EarthRadar UXO Testing at Fort A.P. Hill
Unclassified

5a. CONTRACT NUMBER
DASW01-98-C-0067

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Tuley, M. T. ;
Ralston, J. M. ;
Rotondo, F. S. ;
Andrews, A. M. ;
Rosen, E. M. ;

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME AND ADDRESS
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA22311-1772

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME AND ADDRESS
ODUSD(S&T)
Suite 9030
1777 N. Kent Street
Rosslyn, VA22209

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
APUBLIC RELEASE
,

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This report provides an analysis of EarthRadar blind testing carried out at the Joint Unexploded Ordnance Coordination Office (JUXOCO) Pilot Site, Fort A.P. Hill, Va., during the fall of 2000 and spring of 2001. The EarthRadar, developed by Bakhtar Associates (Newport Beach, Calif.) under an Air Force Small Business Innovative Research contract, is intended to detect and classify buried objects. The report describes the JUXOCO test facility; the EarthRadar hardware, signal processing, and data analysis; and the blind lane test results. It evaluates EarthRadar performance in the context of documented requirements for systems employed in UXO detection/discrimination and draws conclusions regarding the limitations of this technology for UXO clearance applications.

15. SUBJECT TERMS
ground-penetrating radar; UXO; target ID

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 68	19. NAME OF RESPONSIBLE PERSON Fenster, Lynn lfenster@dtic.mil
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER International Area Code Area Code Telephone Number 703767-9007 DSN 427-9007

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39.18

INSTITUTE FOR DEFENSE ANALYSES

IDA Document D-2625

Evaluation of EarthRadar UXO Testing at Fort A.P. Hill

M. T. Tuley
J. M. Ralston
F. S. Rotondo
A. M. Andrews
E. M. Rosen

PREFACE

This effort was undertaken in response to DDR&E's desire to better understand the operation of the EarthRadar system and its potential capabilities. The EarthRadar was designed, built, and is operated by Bakhtar Associates of Newport Beach, California. The Naval Explosive Ordnance Disposal Technology Division (NAVEODTECHDIV) directed the testing, with the Joint Unexploded Ordnance Coordination Office (JUXOCO) providing their Pilot Test Site at Fort A.P. Hill, Virginia, and support for the test site. The Institute for Defense Analyses was tasked to assist in test planning, test oversight, and test analysis and scoring.

We would like to acknowledge the cooperation of Dr. Khosrow Bakhtar, the developer of the EarthRadar, and his colleague, Ms. Ellen Sagal. Their willingness to answer questions, provide data, and explain their data collection and analysis procedures greatly assisted this effort. The guidance provided by Mr. Christopher O'Donnell, NAVEODTECHDIV, and Mr. Atul Patel, JUXOCO, was key to the successful completion of this effort. We would like to also acknowledge the oversight and support provided by Ms. Teresa Puretz, DDR&E, and Mr. Dennis Van Derlaske, Director, JUXOCO. The assistance of Mr. A. Pope Burr, NAVEODTECHDIV, and Mr. Pete Lacko, NVESD, in electromagnetic compatibility (EMC) test design and testing was extremely valuable in resolving EMC issues surrounding the Fort A.P. Hill data collection. Finally, the excellent test site preparation by Mr. John Moulton and Mr. John Cary of E-OIR Measurements, Inc., and their continuous support during testing are greatly appreciated.

CONTENTS

FIGURES

TABLES

EXECUTIVE SUMMARY

SCOPE

This report provides a description of the EarthRadar system and an analysis of its performance in blind tests carried out at the Joint Unexploded Ordnance Coordination Office (JUXOCO) Pilot Site, Fort A.P. Hill, Virginia, during the fall of 2000 and spring of 2001. The EarthRadar, developed by Bakhtar Associates (Newport Beach, California) under an Air Force Small Business Innovative Research contract, is intended to detect and classify buried objects. Thus, it falls into the general class of what are commonly termed ground-penetration radars (GPR).

TEST DESCRIPTION

For these tests, EarthRadar data were collected in calibration grids and blind test grids with the focus on testing the system's ability to perform target discrimination at known locations. The JUXOCO Pilot Site currently contains 140 calibration squares and 260 blind squares. At the center of each 1-m × 1-m grid square, either a UXO item, a clutter item, or nothing was buried. The buried ordnance included projectiles ranging in size from 20 mm to 155 mm, mortars, Rockeyes, and submunitions. The intent of the test is to measure only the sensor operating characteristics. The site and test procedure are designed to remove variables such as navigation accuracy and site coverage not directly associated with the performance of the technology under test.

DATA

The final report submitted by the EarthRadar contractor (Ref. 1) documents the UXO blind tests and presents the results of the data analysis. It provides declarations of the contents of 112 of the 260 blind squares. The EarthRadar report covers the results of only a subset of the blind squares because of contractor problems that slowed data collection and reduction. The main contributor was self-interference of the EarthRadar transmitter with its on-board Global Positioning System (GPS). Although the site grid provides excellent geolocation, EarthRadar software requires GPS position for its data acquisition and subsequent data processing.

For each square declared to contain a UXO object, Reference 1 provides a number of radar plots, including the three-dimensional (3–D) plot from which the declaration is made. The depth, orientation, and assessed ordnance type are called out. For squares declared to contain clutter objects, an assessed object depth is provided in some cases (20 out of 65). Because confidence levels were not provided for declarations, no information on the receiver operating characteristic (ROC) curve of the sensor could be determined.

CONCLUSIONS

Performance Assessment

Standard ordnance-detection equipment performance-evaluation procedures were applied to the declarations reported in Reference 1. The system test results reflect ordnance detection, discrimination, and depth-measurement capabilities that are similar to those that would be obtained by random chance. This performance falls far below that of current state-of-the-art radar systems.

1. EarthRadar declarations did not successfully distinguish between an empty square and one with any type of emplaced content (UXO or clutter):

    −   75 percent of the empty squares were declared to contain UXO or clutter

    −   73 percent of squares containing UXO or clutter were declared to contain UXO or clutter.

2. In squares containing UXO, EarthRadar declarations that a square contained UXO or clutter or was empty were statistically random:

    −   33 percent of UXO squares were correctly declared UXO

    −   38 percent of UXO squares were misidentified as clutter

    −   29 percent of UXO squares were declared empty

    These percentages are consistent with the most probable random values, given the sample size.

3. The errors on declared depths for squares containing emplaced objects (UXO or clutter) are consistent with uniformly distributed random depth guesses.

System Operation Assessment

After observing system operations and analyzing the test data, we identified several causes of the poor performance. Some of these problems are specific to the EarthRadar, but others are endemic to all GPRs:

1. The system concept is fundamentally flawed. For detection, the EarthRadar depends on radar scattering mimicking optical scattering. At GPR wavelengths, including those used by the EarthRadar, scattering from buried targets is specular, not diffuse. With the real-aperture processing and bandwidths used, the EarthRadar resolution is not fine enough in any dimension to provide optical-like target images suitable for discrimination, even if the scattering from the targets were diffuse.

2. The imaging software used by the EarthRadar allows arbitrary scaling and thresholding to be applied to the data separately in each image dimension. Such data manipulation can produce misleading 3–D images, particularly in the case where the target configuration is known.

3. Coupling between the transmit and receive antennas and severe system ringing (self-clutter) obscured target responses and made it difficult to separate desired responses from system artifacts.

In short, the EarthRadar provides no operational, technical, functional or performance benefit to the unexploded ordnance detection problem. The system architecture and signal processing employed are well understood in the GPR community. Given the frequency band and bandwidth constraints under which any GPR must operate, the laws of physics do not support production of optical-like images of UXO targets, images on which EarthRadar depends for target detection and identification. There is no reason to expect that modifications to the current radar could provide enhancements that would allow it to perform better than other currently deployed systems.

1. INTRODUCTION

The EarthRadar, developed by Bakhtar Associates of Newport Beach, California, under an Air Force Small Business Innovative Research (SBIR) contract, is designed to detect and classify buried objects. Thus, it falls into the general class of what are commonly termed ground-penetration radars (GPR). The purpose of the effort described in this document is twofold. First, the government desires an evaluation of the performance of the EarthRadar in a blind test against buried unexploded ordnance (UXO). Second, there is a desire on the government's part to gain a greater understanding of the signal processing and data analysis that drive the performance of the radar.

Section 2 of this document focuses on the facility where testing was carried out. Section 3 describes hardware, signal processing, and data analysis employed by the EarthRadar. Section 4 describes the blind lane test results and evaluates EarthRadar performance. The section also reviews and discusses documented requirements for systems employed in UXO detection/discrimination. Results are provided for other sensor modalities used in UXO clearance, and recent radar results produced by other researchers in this area are included. These documented requirements and recent results provide a benchmark against which EarthRadar performance may be compared. Section 5 evaluates EarthRadar signal processing and analyzes electromagnetic compatibility (EMC) problems within the system that hampered data collection. Finally, Section 6 provides conclusions regarding EarthRadar performance in these tests.

2. TEST FACILITY DESCRIPTION

EarthRadar testing was carried out at the Joint UXO Coordination Office (JUXOCO) Pilot Site at Fort A.P. Hill, Virginia. For these tests, data were collected in two 70-m × 1-m calibration lanes and in thirteen 20-m × 1-m blind lanes, as shown in Figure 1. Only the first 20 squares in each blind lane (shaded in the figure) are populated at this time. Each active lane is separated from the next active lane by a 1-m wide lane where nothing is buried. Each of the active lanes is divided into 1-m × 1-m grid squares. Before lane setup, the entire area was graded and surveyed with a metal detector. All detected metallic objects were removed, and the area was smoothed. At the center of each square, an 18-in. diameter hole was dug with an auger. A UXO object, a clutter object, or nothing was placed in the hole, and then the hole was refilled. The entire area was subsequently rolled to provide a smooth surface. Four plastic pegs, one at each corner, mark each grid square. The burial process creates a somewhat artificial soil context, but it is the same for all cases; that is, digging a hole in each square, whether a target was to be buried in it or not, provides no intrinsic clue about whether a target is buried there. In addition, the digging allows removal of unintended clutter objects from the test area.

The Pilot Site provides a highly controlled site for UXO testing, one where factors such as site coverage and geolocation ability that can affect primary detection ability have been purposely eliminated. Thus, the site can be thought of as primarily testing the ability to do target discrimination at known locations.

For each of the 140 calibration squares, the EarthRadar contractor was provided precise information concerning the identity, depth, and orientation of the buried object (in each square that contained an emplaced object). Within the 260 blind lane squares, the contractor knew that objects of interest occur only in the center of each grid square, but did not know what might be placed in any given square. The contractor has been provided pictures and dimensions of each target type and a range of burial depths for each type target. Small targets tend to be found at real-world UXO sites at shallower burial depths than do larger targets, and that is reflected in the Pilot Site burial depths. Data provided by the U.S. Army Engineering and Support Center in Huntsville, Alabama, were employed to select the range of target depths used (Ref. 2). The following 11 types of UXO were buried in the calibration and blind lanes:

Figure 1. Fort A.P. Hill JUXOCO Pilot Site Layout

- 20-mm, 40-mm, 57-mm, 105-mm HEAT, and 155-mm projectiles;

- 60-mm and 81-mm mortar rounds;

- Mk 118 Rockeye;

- M42 submunition;

- BLU–26; and

- BLU–63.

Clutter items covered a range of sizes and physical configurations, but in each case, scrap items were chosen that did not physically resemble any of the buried ordnance.

The testing occurred in two phases. In the first (4 October 2000–17 October 2000), the contractor collected data in the calibration lanes. The purpose of that collection was (1) to allow him to ensure that his system was operating correctly in the Fort A.P. Hill environment; (2) to allow selection of the appropriate antennas and frequency ranges of operation for the radar; and (3) to allow data collection on known targets of the same types that would be encountered in the blind lanes, which would assist in target identification. The second phase of testing, data collection in the blind lanes, involved several visits of the contractor to the test site over the period from 7 November 2000 through 15 March 2001. The extended testing period led to significant variations in site conditions during data collection. In general, however, the ground was dry for the data collections in 2000 and wetter for the collections in 2001. Data were never collected when the site was extremely wet because of damage to the lanes that would have been caused by the EarthRadar tractor tires and the antenna sled under those conditions.

Apparent radio frequency interference (RFI) causing frequent loss of lock by the EarthRadar differential Global Positioning System (GPS) slowed data collection, thereby significantly extending the planned schedule. Continuous operation of the GPS is required to provide position information for later image formation in the EarthRadar signal processing. Investigations by the government team revealed that the RFI was self-induced, caused by harmonics of the transmitted signal interfering with the local GPS receiver on the sled carrying the EarthRadar antennas. An analysis of the RFI problem is provided in a later section of the report.

3. EarthRadar SYSTEM AND SIGNAL PROCESSING DESCRIPTION

This section provides a detailed description of EarthRadar hardware, data-collection procedures, and signal processing. Each of the major components of the EarthRadar is described, and its function within the system is discussed. Data-collection procedures are then covered, along with what data are recorded. Finally, a detailed description of the data processing is provided, along with some of the implications for choices made in the processing.

3.1 EarthRadar HARDWARE

The EarthRadar is a stepped-frequency GPR built around a Hewlett-Packard (now Agilent Technologies) 8753D vector network analyzer. Additional radar subsystem equipment includes a pair of 5-m long coaxial feed cables and a pair of antennas, one operating as the transmit antenna and one as the receiving antenna. Ancillary equipment includes a differential GPS system [a GPS master station, the mobile GPS system located with the radar, and a very high frequency (VHF)/ultra high frequency (UHF) communications system connecting the two], a GPS data logger, a laptop computer for system control and data display, and a 12-volt battery and inverter to provide alternating current (AC) power to the network analyzer and computer. The network analyzer, mobile GPS equipment, computer, and battery/inverter are mounted on a garden tractor, which drags the antennas, mounted on a plastic sled, behind it. A description of the major system components is provided below.

3.1.1 Network Analyzer

The 8753D vector network analyzer is a standard Hewlett-Packard product that is widely used in the electronics industry to measure the complex scattering matrix parameters (S-parameters) of general two-port networks (Ref. 3). Only two of the four S-parameters (S_{11} and S_{21}) are used in EarthRadar operation. The 8753D provides many features that are not required or used by the EarthRadar, and we limit the discussion here to the features important to this task.

For our purposes, a vector network analyzer, although much more complicated in practice, can be treated as a low-power transmitter and three coherent, homodyne receivers (a reference receiver and one for each of the two measurement ports). In serving as the EarthRadar transmitter/receiver, a calibrated signal is transmitted from port one of the network analyzer into the EarthRadar transmit antenna and received, via the receive antenna, by port two. The reference receiver measures the amplitude and phase of the transmitted signal, which is used to normalize the S-parameters.

In assessing the frequency band over which a given set of antennas may be employed, the contractor takes advantage of the network analyzer's measurement of S_{11}, which is often called the reflection coefficient or the return loss. S_{11} is a measure of the input voltage reflected back into port one (the transmitting port of the network analyzer) due to impedance mismatch between the network analyzer and the transmit antenna and its feed network. It is defined as the ratio of the reflected signal and the transmitted (reference) signal.

In observations of the contractor's selection of antenna operating frequency bands during calibration lane testing, it appears that the band is chosen so that a return loss value of at least 6 dB is achieved (6-dB to 10-dB values appear to be typical). These values indicate that 75 to 90 percent of the available power from the network analyzer is being accepted by the transmit antenna system. For many systems these would be marginally acceptable efficiencies; however, for GPR, detection is most often clutter rather than noise limited, and so achieving very high efficiency is not usually a major concern. Of more concern is the fact that this level of mismatch will cause "ring down" in the system (i.e., produce self-clutter) that can hide the signal response from targets during the period in which the ringing decays. In fact, based on our analysis of the test data, self-clutter proves to be a major limiting factor in EarthRadar performance. The return loss characteristics are not recorded by the contractor and are simply employed to select a usable operating band for a given antenna set at a given measurement site.

The more important scattering parameter, and the one recorded for further processing to determine the presence of a signal response and for target detection/ identification, is S_{21}. S_{21} (also called the transmission coefficient) is the signal received by the second port of the network analyzer (connected to the receive antenna) due to excitation of the network by the first port. Analogous to S_{11}, S_{21} is the port 2 receiver response divided by the transmitted signal (i.e., the reference receiver response).

For operation across a desired frequency band, the network analyzer generates a coherent, linear, frequency-modulated (FM) waveform, often called a "linear chirp," that begins at a frequency lower than the chosen frequency band and continues past the highest desired frequency. The network analyzer samples a number of evenly spaced, discrete frequencies over the selected band, as chosen by the operator. For example, in the blind lane data collection, 201 frequencies were sampled, typically over an 800-MHz to 1-GHz chirp (i.e., samples were collected in 1-MHz steps). The network analyzer automatically adjusts the chirp rate to match the operator-selected receiver bandwidth (300 Hz was used in this testing). A coherently transmitted and received stepped-frequency signal can be processed through an inverse fast Fourier transform (IFFT) to provide improved range resolution over that available from an unmodulated, narrowband, pulsed waveform. In general, such stepped-frequency processing can be assumed to provide, at best, a range resolution $\Delta R = c/2B$, where c is the speed of light in the medium and B is the bandwidth (Ref. 4). One disadvantage of using a stepped-frequency waveform is that it is ambiguous in range. That is, if the propagation range of the signal is large enough, targets at long ranges will fold down and appear at shorter ranges. This can create additional clutter in the area of the target ranges of interest and provide incorrect target location information. The 1-MHz frequency step in the example above, however, results in a free-space unambiguous range interval [$c/2\Delta f$, where Δf = the frequency step (Ref. 4)] of 150 m and ensures that range-aliasing will not occur, even in soils with high indices of refraction.

While the network analyzer employs a relatively complicated triple-conversion receiver, the effect is to mix the received signal down to direct current (DC) using the transmitted signal as a reference. Quadrature mixers provide in-phase (I) and quadrature-phase (Q) measurements at each frequency that are then digitally sampled. Amplitude ($[I^2+Q^2]^{1/2}$) and phase ($\tan^{-1}[Q/I]$) can be derived from the digital I and Q samples. The S_{21} I and Q data are recorded for processing in the IFFT.

3.1.2   Antennas

Although details concerning the construction of the EarthRadar antennas are not available (the antennas and feed networks are encased in a sealed composite enclosure), the developer has provided some general information concerning antenna designs. Two different antenna types were used in the data collection. The developer generally describes one type as "center-fed dipoles." Figure 2 provides a photograph of a set of the

Figure 2. Photograph of Typical EarthRadar Antennas

dipole antennas on the sled used to drag them. The dipoles are fed by a lumped-component feed network used to provide impedance matching in the operating environment. As noted in the section above, a frequency band over which the antenna system exhibits a 6- to 10-dB return loss is typically chosen. The antennas are parallel to each other and have parallel polarizations, which result in significant direct coupling between the antennas. Direct-coupling effects, which are always apparent in the data displays provided by the radar, are discussed later.

The dipole antennas and sled ride on, or nearly on, the soil surface. This is typical for GPR dipole antennas, and close contact is usually chosen to provide efficient coupling of the radar signal into the ground (i.e., to reduce the large reflective loss due to the impedance discontinuity represented by the air/soil boundary). Unfortunately, this also causes the ground to act as an integral part of the antenna system, thereby making antenna characteristics vary with soil conditions. Such effects are endemic to the application of this antenna technology and can produce varying and unpredictable performance in a survey, often within very short distances.

One unusual characteristic of the EarthRadar dipole antennas is their orientation. Most GPR equipment (whether impulse or step-frequency) orient dipole antennas with the axis cross-track. That orientation places the broadest part of the antenna pattern along-track, thus aiding the formation of the familiar hyperbolic shape that represents a potential target response. Because the EarthRadar does not use recognition of a hyperbolic pattern in its detection processing, however, the antenna orientation in this case is likely not critical.

For higher frequency operation in the calibration lanes, a pair of "horn" antennas mounted in the same type enclosure was employed. Again, no details are available, but the developer indicated that a lumped-circuit matching network similar to that with the dipoles is employed and that the two horns also have parallel polarizations.

3.1.3 GPS Equipment

The EarthRadar uses differential GPS to provide relative position coordinates (Northing, Easting, and Elevation) for each of the data-collection events. A data-collection event, termed a "trace" by the contractor, comprises a sample at each of the frequencies over the band being used (e.g., one I and one Q sample at each of 201 frequencies). Tractor motion is typically slow (<0.1 mph), stepping through the entire frequency band is done in about a half-second, and the coordinates are assumed to be constant over the time required to collect a trace. Note, however, that the antenna can move about 2 cm during the collection period, which is of the order of the ±1-cm differential GPS accuracy stated by the developer (Ref. 5). In addition, the position sensed is actually the position of the GPS antenna on a 7-ft tall support pole (visible in Figure 2) mounted to the sled. The elevation difference between the EarthRadar and GPS antennas is removed in processing. Test observers noted significant side-to-side motion of the GPS antenna with sled movement, and the position variations were evident in several of the EarthRadar data files analyzed by IDA. Thus, actual position data are likely much less accurate than advertised by the radar developer.

A master GPS station is set up on the site near the UXO lanes to provide a good view of the sky for maximum satellite reception. A VHF/UHF system with several choices of frequency bands is used to communicate between the two GPS systems and provide differential GPS. During calibration lane testing and early blind lane testing, significant problems were encountered with loss of the GPS signal. The radar developer initially thought that the problem was interference with the radio link between the master and slave GPS systems. The Naval Explosive Ordnance Disposal (EOD) Technology

Division provided personnel and a Hewlett-Packard 8562A spectrum analyzer to evaluate the environment. Although a number of signals were found within 1 MHz of the GPS communications signal, none was within 20 kHz of the frequency being used by the EarthRadar system, and all signals were at least 25 dB below the desired signal. Later, analysis by the Institute for Defense Analyses (IDA) raised the likelihood that the interference was caused by harmonics of the network analyzer transmit signal interfering with the local GPS. That hypothesis was confirmed, and the EarthRadar operator chose a frequency band for blind lane collection (800 MHz–1 GHz) where the second and higher harmonics fell above the GPS L1 frequency of 1,575.42 MHz.

3.2 DATA COLLECTION

The EarthRadar collection and display parameters are controlled by a graphical user interface (GUI), written using the Lab Windows CVI commercial software package, and appearing on an input screen of the control and data-acquisition computer. In describing the data-collection process and the displays that are produced, we have included some specific examples taken from Reference 1, the final report on the blind lane testing.

Figure 3 provides an example EarthRadar set-up GUI. Several of the parameters displayed are important to later discussions. Note that the frequency range covered is 800 MHz to 1,000 MHz, a total bandwidth of 200 MHz. The system uses 201 frequency steps, so each frequency step is 1 MHz. The IFFT array size is 1,024, and a Hann window is used. This indicates that in subsequent processing for data display, the 201 complex data points are windowed and then padded with zeros to a total data vector length of 1,024 before inverse transforming. A wave speed of 8.5 cm/ns is used, implying a soil index of refraction of 3.53 and a relative permittivity around 12.5. For the JUXOCO Pilot Site, measured in situ relative permittivities were in the range of 7.5 to 11.1 in the frequency band of interest (Ref. 6). IDA calculated a wave speed of 9.9 cm/ns, based on calibration lane data collected for two flat disc targets, one buried 1 ft deep and the other 2 ft deep. That corresponds to a relative permittivity of 9.13, which is more in line with the values of Reference 6. However, even if the wave speed were constant over the entire site and the IDA calculated value were correct, use of the 8.5 cm/ns value would result in less than a 15-percent depth error, which is within the depth resolution of the system for the target depths of interest. Therefore, it should not constitute a significant source of error.

Figure 3. Typical EarthRadar Setup Screen (Ref. 1)

After input parameters are entered via the input GUI, EarthRadar data are collected by triggering the network analyzer to collect a set of I and Q data points, one at each frequency. The data acquisition computer receives a trigger signal from the GPS data logger indicating that differential GPS lock is present and ordering a position point to be recorded. The data logger signal simultaneously triggers the network analyzer to record a set of frequency points. I and Q data from each set of frequency points (i.e., each trace) are stored in a separate file on the data-acquisition computer. The data acquisition computer takes the I and Q file, windows the data using a Hann window, pads the data string to 1,024 points, computes an IFFT, and displays the real part of the IFFT (subsequently indicated as Re{IFFT}) on the computer screen. After display, the computer waits for the next trigger from the GPS system (nominally set for approximately one data point each second), at which point the data collection is repeated.

The data-collection procedure begins with the EarthRadar stationary. Approximately 25 traces of data are normally collected before the radar moves. At the standard tractor speed in the calibration lanes, one data trace is collected every 1 to 2 in. of antenna travel. Data collection continues while the tractor is moving, and then another 25 or so traces are collected at the end of the run, after the tractor stops.

A manual log that notes the trace number at which motion started, the trace number at the center of each calibration square, and the trace number at which motion stopped was kept during the data collections. Thus, in subsequent data analysis, the operator knew the trace numbers encompassing the center of a grid square in which to look for a target response. The software used to provide a three-dimensional (3–D) image of the target requires at least two cuts in different directions over the target position. For the JUXOCO Pilot Site lanes, the tractor pulled the antennas along the length of the lanes to produce one set of data. A second cut was obtained with the tractor pulling the antennas across the lanes. For some targets, diagonal cuts were also collected.

The data used by Bakhtar Associates to determine whether a signal response is present are the real part of the IFFT from the traces, which can be displayed in a number of processed waterfall plot formats. A 3–D model is created for imaging by merging the Re{IFFT} points with GPS data, attributing a set of coordinates to each IFFT point (Northing, Easting, and Elevation) using an assumed speed of wave propagation, and providing this 3–D data file as input to the RockWare, Inc., commercial software package that creates the image model. Details of the data processing and presentation follow.

3.3 DATA PROCESSING AND PRESENTATION

Data processing in the EarthRadar is relatively simple and straightforward. As noted earlier, the fundamental data product on which detection decisions are based is the collection of traces of the Re{IFFT} produced from the complex stepped-frequency data. There are a number of subtleties, however, in the results of the signal processing performed that must be appreciated if system operation is to be truly understood. The descriptions of data processing and presentation provided here are based on published information about the EarthRadar (e.g., Ref. 5) and on discussions with the EarthRadar developers. In particular, the government and IDA traveled to Bakhtar Associates' facility (Newport Beach, California) in October 2000 to see a demonstration of data processing on calibration-lane data. Questions during that demonstration were used to confirm our understanding of the processing.

In discussing the signal processing, we first consider the simple case of a point target in a homogeneous medium and determine the characteristics of the target return after EarthRadar-type processing. This represents a much simpler case than that faced by the radar, where reflections from multiple targets occur and the earth may not be homogeneous, but it does give an idea of the limits in an ideal case.

Assume a point target buried at a depth D meters below the radar antennas on the surface. Without loss of generality, we can assume the amplitude of the return at each of the step frequencies is unity, so the signal out of the network analyzer at the kth frequency step is given by

$$F(k) = \exp\left(-j\frac{4\pi(f_0 + k\Delta f)D}{c}\right),\ 0 \le k \le N_0 - 1$$

for the N_0 frequency steps, and

$$F(k) = 0,\ N_0 \le k \le N - 1 \tag{1}$$

for the zeroes padded at the end of the data vector, where $j = \sqrt{-1}$, f_0 is the initial (lowest) frequency, Δf is the frequency step, and c is the wave speed in the medium. The inverse discrete Fourier transform (IDFT) is typically defined as

$$f(n) = \frac{1}{N}\sum_{k=0}^{N-1} F(k)\exp\left(j\frac{2\pi nk}{N}\right)\qquad n = 0...N-1\ , \tag{2}$$

where each element of $f(n)$ represents a sample at a time $n\Delta t$, and $\Delta t = 1/(N\Delta f)$.

After inserting $F(k)$ in the IDFT formula and rearranging terms, we find that

$$f(n) = \frac{1}{N}\exp-\left(j\frac{4\pi f_0 D}{c}\right)\sum_{k=0}^{N_0-1}\exp\left(-j\frac{4\pi k\Delta f D}{c}\right)\exp\left(j\frac{2\pi nk}{N}\right)\ . \tag{3}$$

Carrying out the appropriate algebra, the return from a point target can be expressed as

$$f(n) = \frac{1}{N}\exp\left(-j\frac{4\pi f_0 D}{c}\right)\exp\left[j2\pi\left(\frac{N_0}{N}\frac{n}{2} - \frac{\Delta f D N_0}{c}\right)\right]\frac{\sin\left[\frac{N_0}{2}\left(\frac{2\pi n}{N} - \frac{4\pi\Delta f D}{c}\right)\right]}{\sin\left[\frac{1}{2}\left(\frac{2\pi n}{N} - \frac{4\pi\Delta f D}{c}\right)\right]}\ . \tag{4}$$

Taking each of the factors of Equation (4) in order, the $1/N$ factor is simply a normalizing factor resulting from the particular definition of the IDFT being used. The first exponential factor represents a constant phase shift applied to each term of the IDFT. The specific value of the phase shift depends on the initial frequency and the depth of the target. The second exponential factor is a sampled complex sinusoid whose frequency is $N_0\Delta f/2 = (B+\Delta f)/2 \approx B/2$, where B is the bandwidth of the stepped-frequency waveform. Note that the frequency is not exactly $B/2$. The difference arises because the DFT assumes that it samples a periodic function. For that reason, there is an implicit sample point beyond the final actual sample point that makes the effective bandwidth $B+\Delta f$,

rather than B (Ref. 7). Finally, the $\sin(0.5N_0\theta)/\sin(0.5\theta)$ factor, called the Dirichlet kernel (Ref. 8), is the impulse response of a sampled rectangular spectrum with bandwidth $N_0\Delta f \approx B$, offset so the peak occurs at the target depth, D. Applying a window function, such as the Hann window, leaves the form of the IDFT the same, except the Dirichlet kernel is replaced by one appropriate to the window function.

For an unpadded IDFT, $N_0 = N$, and the complex sinusoid in the second factor of Equation (4) will be sampled at points 180 deg apart, with the location of the sample points controlled by the $\Delta fDN_0/c$ phase term. This results in a sawtooth waveform (i.e., alternately positive and negative) for the real or the imaginary part of the IDFT. Padding by adding zeroes to the end of the data vector per Equation (1) moves the sampling interval closer together by a factor of N_0/N, but does not change the frequency of the complex sinusoid.

That a complex sinusoid appears at a frequency of $B/2$ is inherently an artifact of the form of the IDFT and the manner in which padding is applied. It has nothing to do with the physics of the problem, but instead depends on the details of the processing, as explained below.

After insertion of the point-target return into the IDFT formula, shown in Equation (3), we find that the starting transmitted frequency (f_0) appears only in the constant phase-shift factor. The summation in Equation (3) operates on two factors: the function to be transformed and the kernel of the IDFT. Within the function to be transformed, the frequency is represented by the $k\Delta f$ product. Because we have moved the starting frequency outside the summation and the computer algorithm defines k to run from 0 to $N_0 - 1$, this, in effect, puts the first frequency in the summation at 0 Hz and the final frequency at $(N_0 - 1)\Delta f = B$ Hz. The result is to generate a sampled complex sinusoid with a frequency $\approx B/2$ Hz, as is seen in Equation (4). We can think of this as an effective carrier frequency waveform multiplied by the window function around the target location. If, instead, the algorithm were written so that $-N_0/2 \leq k < N_0/2$, the frequency band would be symmetric around zero frequency (i.e., run from $-B/2 \leq f \leq B/2$), and there would be no effective carrier (i.e., its frequency would be 0 Hz).

This focus on the details of the IFFT processing used by EarthRadar has two goals. The first is to explore the expected form of the data as an aid to understanding later results. The second is to point out that for the data employed by the EarthRadar, the Re{IFFT} does not represent a signal that is physically meaningful. The effective carrier is an artifact of the IDFT definition used, and the response from the target represented by

it is a function of the range to the target. By changing the IDFT definition or how the zero padding is applied, we can change the appearance of Re{IFFT}.

Figure 4 provides a representation of the response for the 800 MHz–1GHz waveform used for the blind lane collection for two unit amplitude point targets, one at a depth of 0.2 m and the other at a depth of 1.0 m. A relative permittivity of 8.94 and a conductivity of 0.013 mho/m were used to match measured JUXOCO Pilot Lane properties at 910 MHz for dry soil (Ref. 6). An IDA Mathcad model that processes I & Q samples in an identical manner to the EarthRadar generated the synthetic data. That is, the return (amplitude and phase) from the point targets was computed for 201 frequencies evenly spaced over the 800–1,000 MHz frequency band, a Hann window was applied to the complex frequency-domain data, and the windowed 201 points were padded at the end with 823 zeroes to provide 1,024 points into an IFFT routine. Figure 4, however, presents only a small subset of the 1,024 data points derived from the IFFT. Only the interesting region around the input ranges of the point targets (0–1.6 m) has been plotted because the values of the remaining IFFT points are very close to zero. The dotted curve is the magnitude (envelope) of the IFFT, normalized to its maximum value. The solid curve is the real part of the IFFT (the data used by the EarthRadar), also normalized to the maximum value of the IFFT magnitude. Because this simulation implicitly assumes that the antennas add no distortion to the signal, it represents an ideal impulse response for the chosen bandwidth.

The shape of the dotted curve showing the signal envelope is independent of the details of the IFFT algorithm employed to process the I and Q data and of the ranges of the point targets. Note that the 4-dB width (distance between points where the amplitude is greater than 0.63) of the signal envelope is about 0.4 m, corresponding to the 0.24 m resolution for a rectangular spectrum [$c/(2B*\sqrt{\varepsilon_r})$], broadened 49 percent by the Hann window (Ref. 7).

The specifics of the real part of the IFFT, however, are very dependent on the IFFT algorithm, as discussed above, and also on the range because of the constant-phase factor $\exp(-j4\pi f_0 D/c)$ that appears in the IDFT. The real part of the complex sinusoid represented by the second exponential factor in Equation (4) peaks near the target. That the constant phase-shift term is a function of the target range, however, modifies the phase so that the real part of the product of the terms does not necessarily peak at the target. We chose the depths shown in Figure 4 to illustrate that point. For Target 1, the windowed sinusoid shows a positive peak somewhat shallower than the target location

Figure 4. Simulation of the Response of Two Point Targets Buried at 0.2 m and 1.0 m for a 201-Point Stepped-Frequency Waveform Covering 800 to 1,000 MHz, Hann Windowed, and Zero Padded to 1,024 Points. Soil dielectric properties are typical of those measured at the JUXOCO Pilot Site for dry soil at 910 MHz.

and a negative peak somewhat deeper. In contrast, the target location for the second target is near the negative peak of the sinusoid. As expected, we see that the wavelength of the sinusoid is about 0.5 m, corresponding to a frequency around 100 MHz (B/2), after accounting for the two-way range and the relative permittivity of the soil.

Also note the reduced response from the deeper target (Target 2). This illustrates a typical problem faced by GPRs. Wide bandwidths are generally desired for finer range resolution, which provides better target identification capabilities, but the higher frequencies are attenuated more rapidly with depth, making them unsuitable for deeply buried targets. In fact, dry soil properties were used in creating Figure 4, because typical conductivities (0.039 mho/m vice 0.013 mho/m) attenuate the signal so rapidly that the deeper target is almost undetectable on a linear scale.

Because the form of the response from a target is a function of its range, it would be very difficult to design algorithms that might automatically recognize a target response based on the Re{IFFT} representation. On the other hand, the magnitude of the IFFT does represent a signal that is physically based, is the same regardless of target range, and only depends on the bandwidth of the stepped-frequency signal, as shown in Figure 4.

That is the reason it is much more common in signal processing to work with the magnitude than either the real or imaginary part of the IFFT alone.

3.4 DISPLAY PROCESSING AND FORMATS

The EarthRadar processes the real part of the IFFT in a number of ways so that the operator can visually determine whether a signal response is present. The three major formats used are the wiggle plot, the color plot, and the threshold plot. Other plotting options are discussed in Reference 5 but are not covered here. One attribute common to all of the plots is that the amplitudes of the data are typically modified before display. That process is described first, before the specific display formats are covered.

3.4.1 Linear Contrast Expansion

In displaying the Re{IFFT}, large positive and negative swings will be seen in the initial portion of the trace due to direct coupling of the transmit and receive antennas. If the available gray-scale or color palette is linearly mapped over the entire dynamic range of the signals, subtle amplitude changes indicating the potential presence of a target or clutter scatterer may map to the same value as slightly differently valued surrounding pixels.

To provide contrast within the range of amplitudes expected of targets, a linear contrast expansion is applied to the data before further processing. This contrast expansion combines with a fixed map relating amplitude to color palette (or gray-scale) to clip large positive or negative amplitude values and amplify small amplitude changes among small and intermediate values of the output. Re{IFFT} data are rescaled by dividing each data value by a selectable constant (typically in the range of 0.1 to 0.2). This, in effect, clips the data by forcing all large positive amplitudes to a particular color or shade of gray and all large negative amplitudes to a different color or shade of gray. Small amplitudes are stretched to fill the dynamic range between the extremes.

The linear-contrast-expansion value chosen is a function of the particular plot type designated and of operator preferences and experience. Often, in looking for the presence of a signal, the EarthRadar operator has been observed to try several linear-contrast-expansion factors before settling on one for a particular target location that gives the "best" results.

When two sets of traces taken in different directions are combined for the 3–D modeler, a linear contrast expansion is applied separately to the two data sets. The values

chosen can be different for the two sets. The potential impact of linear expansions on 3–D images is discussed further in the discussion of the EarthRadar imaging capability.

3.4.2 Wiggle Plots

The developer calls the most basic plot employed by the EarthRadar a wiggle plot. This general type of plot is more commonly called a range-time-intensity (RTI) plot. A wiggle plot, illustrated in Figure 5, simply produces a waterfall plot built up of a succession of traces after linear contrast expansion of the data. Positive values in the plot have the area between the curve and the axis shaded white and negative values have a black background shade. Although related to antenna coupling rather than target return, the shape at the top of the plot echoes the negative/positive/negative signal swings seen for Target 1 in Figure 4.
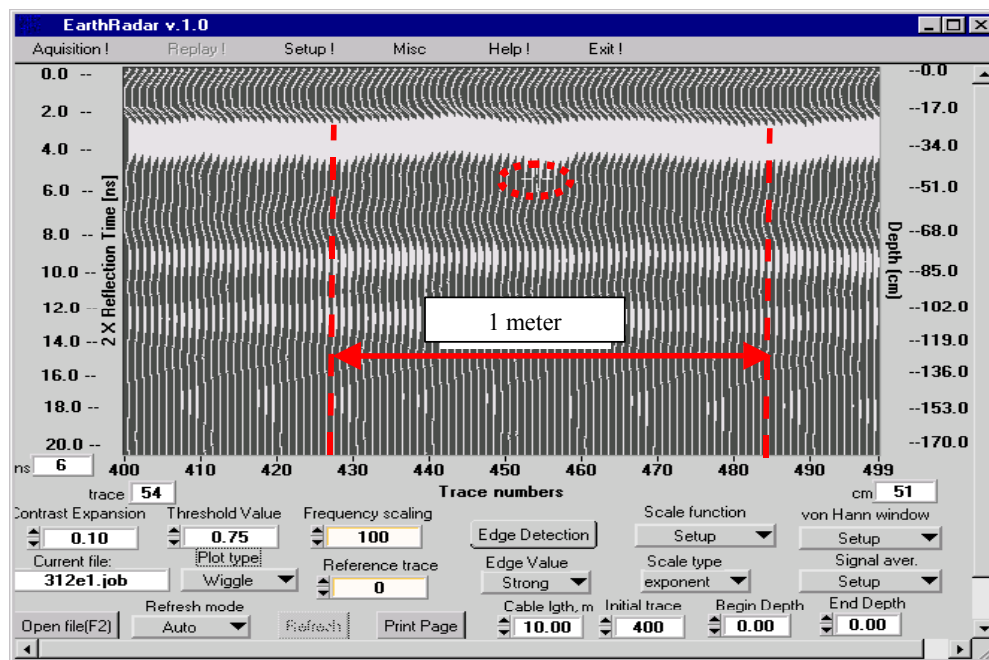
Figure 5. EarthRadar Wiggle Plot (Ref. 1)

Although some minor perturbations of the return can be seen at various depths, this plot is not normally used to determine whether a signal response is present. Instead, the EarthRadar operator uses it to determine the reference time delay for depth calculations. Based on observation and the government team's questioning of the EarthRadar contractor, the bottom of the strong white band caused by direct antenna coupling is used as the reference time from which depth is then determined. It is not obvious why the lower edge of the strong positive trace is chosen, rather than some other point. Although such a choice could bias the estimate of apparent depth, whether it did so could not be

3-14

determined in this analysis because of the lack of correlation between actual and assessed target depths.

Both time and depth scales are provided on this and the other various plots. The time scale on the left side of the plot is indicated as the two-way transit time (i.e., 2 × Reflection Time on the figure), and the transit time from the network analyzer through the transmit and receive cables has been subtracted out. Because the depth scale on the right is obtained by multiplying the time scale on the left by the assumed wave speed, however, the time on the left scale must be the one-way transit time. The right scale is the depth in centimeters for a wave velocity input by the operator (e.g., for the example plot shown, the assumed velocity is 8.5 cm/ns). Generally, the EarthRadar operator determines wave velocity by incrementally separating the transmit and receive antennas and noting the change in time delay as a function of change in distance. That technique has the disadvantages of being dependent on only the near-surface dielectric properties and dominated by direct coupling.

3.4.3   Color Plot

The plot generally used to determine the presence or absence of a signal and the depth of the response is called a color plot by the EarthRadar developer. Figure 6 provides an example of a color plot, again taken from Reference 1. This plot uses a linear contrast expansion of 0.1, and thus each Re{IFFT} value has been multiplied by 10. In addition, interpolation has been used to create a smooth plot. Interpolation is employed in both directions (in the depth direction and across traces), and an inverse distance weighting from neighboring points is employed to select the color plotted. The effect of the linear expansion is to clip the peaks of the direct-coupling response, providing the very abrupt red (maximum negative), black (maximum positive), red pattern. Smaller amplitude variations are then emphasized, as can be seen from the plot. The developer identifies a target response as the pattern around trace 55 at approximately 5.5 ns on the time scale.

An obvious question is why the particular pattern anomaly in the plot was chosen as the target instead of other similar looking areas. In this case, the location of the target was known to be at the center of the 1-m square, and the other responses occurred at different locations. Thus, knowledge of exact target location provided a cue to the radar operator that might not be available, or might not be as accurate, in more realistic testing.
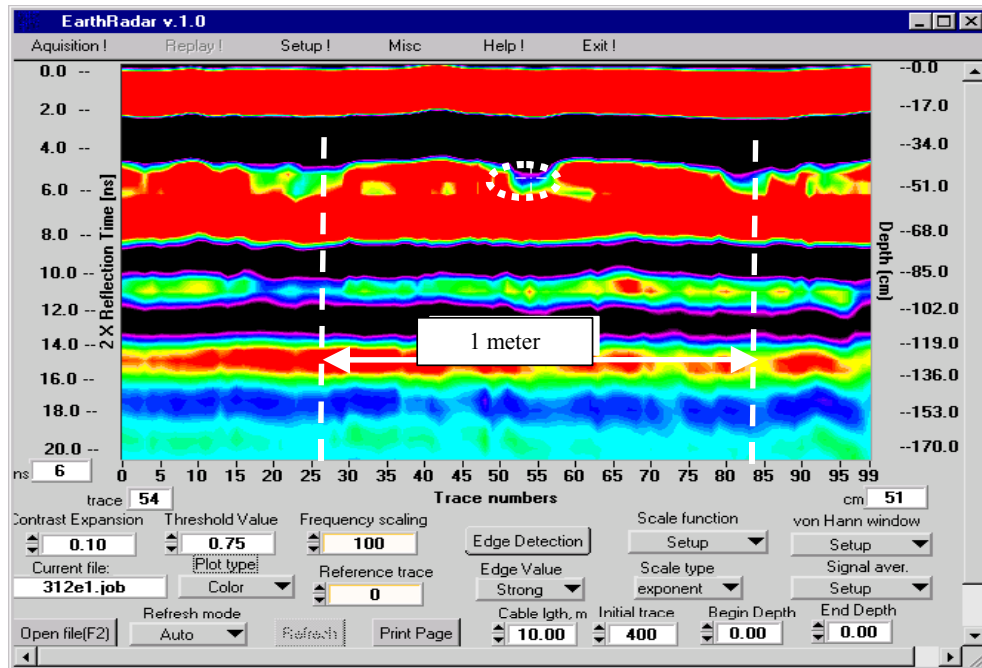
Figure 6. EarthRadar Color Plot (Ref. 1)

Under normal circumstances, if a determination is made that a signal is present at a given location, a second set of data is taken in a different direction to confirm the location and to provide data for later 3–D image generation. Often the two cuts will be at right angles to each other, but they need not be.

For these tests, cuts in two orthogonal directions were made for each grid square to ensure that data were available for 3–D image production on those grid squares where the presence of a target or clutter was suspected. Much of the data processing was accomplished at the Bakhtar Associates' office after data collection was complete.

3.4.4 Other Plots

Three other plots are available that, like the wiggle plot and color plot, are also formed off the RTI data. These are a gray-scale plot, a threshold plot, and an edge-detection plot. The gray-scale plot simply maps the smoothed amplitudes into a gray-scale rather than color palette. The threshold plot, shown in Figure 7, is a black and white plot, where all pixels above the threshold map as white and all pixels below the threshold map as black. Note that because the Re{IFFT} data are bipolar, thresholding suppresses large negative returns (such as the return from Target 2 in Figure 4). Also note that the threshold plot emphasizes the repeated reflections between the transmit and receive

3-16
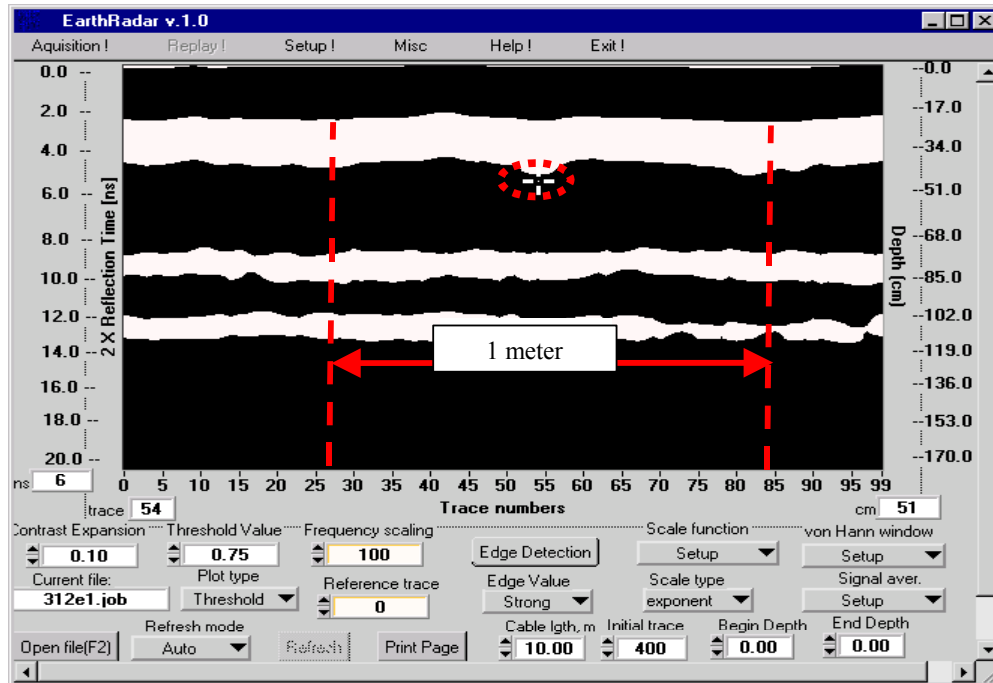
Figure 7. EarthRadar Threshold Plot (Ref. 1)

antennas that produce system self-clutter, as illustrated by the bright, repeated, horizontal bands across the display. Finally, the edge-detection plot employs a standard edge-detection algorithm to create a black and white plot, where edges appear in white. Although the gray-scale and edge-detection plot formats are available, it appears from observation of testing and from the final report that signal presence decisions are principally made using the wiggle, color, and threshold plots.

3.5 3–D IMAGE PRODUCTION

Although the color plot is used to decide whether a signal response of interest may be present, the resolution is not sufficient for detection (decision whether a target signal rather than a clutter signal is present) or classification purposes. For those portions of the process, a 3–D image is created. Figure 8, which illustrates such an image, is used to explain the processing.

As noted above, 3–D image models are produced using the Re{IFFT} data from cuts across the target in at least two different directions. GPS data are required for each trace to calculate Northing, Easting, and Elevation positions for each Re{IFFT} point. A commercial software package developed by RockWare, Inc., creates the 3–D model and

3-17

Figure 8. EarthRadar 3-D Image Plot (Ref. 1)

displays the image. The RockWorks 3–D Visual Pro package consists of RockWorks99, which creates the 3–D model, and Slicer Dicer for model display. RockWorks99 expects a 3–D grid of data point amplitudes as inputs from which a 3–D model is created.

For purposes of explanation, let us assume that the data consist of two orthogonal cuts across the target. Based on the GPS position recorded for each trace, local coordinates of the radar antennas are calculated. Ground slope is determined using the positions of nearby traces. Based on the differential GPS position, the slope of the ground, and the assumed wave velocity in the earth, an assumed position can be calculated for each of the Re{IFFT} data points. Note, however, despite this assumption, this set of positions is not unique. In fact, returns in a given bin come from all points whose transmit-receive time delay is appropriate. For this nearly monostatic system, the set of points with equal time delay is represented by a nearly spherical shell. The amplitude of a scatterer within the shell is the product of the scatterer amplitude and the transmit and receive antenna gains in the direction of the scatterer. (Because we do not have measurements of patterns for the EarthRadar antennas, it is difficult to assess the spatial resolution provided by the system. However, because dipole patterns are typically very broad, resolution is likely to be poor.) These points and their associated amplitude after linear expansion then form a sparse 3–D grid. As the cuts cross in the vicinity of the target location, however, the grid is denser in that region.

3-18

RockWorks99 provides a number of options for creating 3–D models. The one selected by Bakhtar Associates is an inverse-distance-squared algorithm. That is, a dense 3–D grid is created by the software, and the value at each point in the grid is the sum of the values at all the data points, each weighted by the square of its inverse distance from the grid point. Dividing by the sum of the squared inverse distances then normalizes the value at each point. If a grid point falls at the location of a data point, the amplitude of the measured data point is preserved (i.e., the inverse-distance algorithm is not applied). Thus, no focusing algorithm is employed to create the 3–D model. In fact, the squared inverse distance weighting can be thought of as a low-pass, smoothing filter. Such a filter degrades, rather than enhances, resolution.

Although the output of the RockWorks99 3–D modeler is called the 3–D model, significant additional processing is involved before the desired image is produced. The Slicer Dicer software used to display the target actually does most of the image production work. Based on demonstrations of data processing by Bakhtar Associates, the model out of RockWorks99 appears as two intersecting strings of clouds. Slicer Dicer is first used to select only the region in 3-space around the presumed target location. The operator then applies different thresholds, viewing angles, and, possibly, linear-expansion factors in each dimension in an attempt to create a target image. The process is very much qualitative and operator-driven. In the demonstrations observed at Bakhtar Associates, the type target and its orientation were known by the operator attempting to produce the image. In the blind tests, where an unknown target was to be imaged, the ability to successfully detect UXO based on this processing was not evident.

4. DATA DESCRIPTION AND RESULTS

This section provides a description of the analysis products delivered to the government by the EarthRadar contractor in his final report and provides detection and false-alarm performance results. The description of the performance is constrained by the understandable desire of the government to maintain valid blind lane status for the site. Thus, results from individual squares are not indicated, but instead, aggregated performance results are provided. To set the stage for those results, the first subsection describes performance requirements. That is followed by a description of the parameters used for the blind lane collection and aggregated analysis results. Finally, those results are translated into performance.

4.1 PERFORMANCE REQUIREMENTS

As is the case for any detection system, unexploded ordnance detectors that achieve a high probability of detection (P_d) with a minimal probability of false alarm (P_{fa}) are required. UXO detection is required in three different mission areas: environmental remediation, EOD, and active range clearance (ARC). These three mission areas have different requirements for sensor performance (Ref. 9).

In the environmental remediation mission, land that is no longer needed by DoD is cleared with a goal of making it safe for civilian use. Depending on future use plans, the required depth of clearance can range from surface only for minimal activities, to depths of a foot or two for agriculture, to depths in excess of 10 ft for construction. For any specific depth requirement, the P_d must be very high. For example, the Kaho'olawe Tier II prove-out requirement is 0.85 (Ref. 10). As regulators and stakeholders take increasing interest in UXO clearance operations, P_ds near 100 percent will be required (Ref. 11).

The cost of cleanups is largely driven by false alarms, with the Army Corps of Engineers reporting that the majority of a typical cleanup's costs go to digging holes that do not contain ordnance. One goal in environmental remediation research and development (R&D) is to reduce the false-alarm rate at high P_d by applying discrimination techniques to sensor data. Regulators generally find unacceptable the prospect of detecting anomalies and, based on a discrimination decision, leaving potential ordnance

in the ground. Thus, attempts to decrease false alarms by discriminating UXO from clutter must retain virtually 100-percent of the ordnance.

The EOD mission involves the removal of immediate threats in emergency responses or UXO clearance in military missions in support of combat operations. Because safe, unrestricted public access is not involved, the P_d requirements for EOD detectors may be more forgiving, but they remain high. In addition, because the mission must be achievable in adverse conditions, other requirements regarding time line and ruggedness are greater (Ref. 9).

Active range clearance, maintaining ranges used by DoD for training and testing, is different still. The primarily focus is on surface clearance, although at times the recovery of special items from testing programs may require the detection of deeply buried items.

4.1.1 Performance of Magnetometers and Electromagnetic Induction (EMI) Sensors

Technologies currently fielded for UXO detection are the magnetometer and EMI sensor. Until recently, all clean-up work was done using a "mag and flag" approach, where an operator listening to audio output of a magnetometer or EMI instrument placed flags in the ground for excavation. This approach has been shown to be inefficient at finding ordnance, with P_ds in controlled demonstrations falling in the 30–70 percent range. Because there is no ability to distinguish the origin of an anomaly, all anomalies must be investigated and false-alarm rates are very high (Refs. 12, 13).

More advanced systems using magnetometers and EMI are in the R&D phase, and some are beginning to be deployed in field work. Those systems collect geo-referenced data for physics-based analysis or template matching. In controlled tests, these systems have shown P_ds of 95 percent or more. For example, at JPG Phase III, the NRL Multisensor Towed Array Detection System (MTADS) detected 97.5 percent of the ordnance with 41 false alarms per hectare (Ref. 14). In an Environmental Security Technology Certification Program (ESTCP)-sponsored demonstration at JPG in 2000, advanced EMI sensors achieved P_d of greater than 0.85 with 70–100 false alarms per hectare, and in some conditions, P_ds of 100 percent were achieved on larger ordnance. At more stressing sites, however, the same sensors have experienced considerably lower P_ds (Ref. 15).

Some success has been demonstrated in removing false alarms using discrimination techniques on magnetometer and EMI data. For example, MTADS EMI data has been used to estimate the polarizability tensor to determine if objects are prolate like ordnance or oblate like a significant fraction of clutter. In controlled tests, about 30 percent of false alarms were eliminated before reductions in P_d were suffered (Ref. 16).

4.1.2   Performance of Other Radars

GPR applied to the UXO problem in a search mode has generally not been successful in field demonstrations. In the early rounds of the JPG demonstrations, radars were included in the mix of sensors tasked to detect, locate, and identify individual ordnance items in an open field. No radar showed any detection capability in JPG Phases I or II (Refs. 17, 18).

Radars in more constrained tests have shown some capability. The ARL BoomSAR has detected large, near-surface ordnance and showed some discrimination ability in tests against regular arrays of targets at Yuma Proving Ground and Eglin Air Force Base. A P_d of about 0.6 was achieved, with 1,000–10,000 false alarms per square kilometer in the test at Yuma (Ref. 19).

In a closer analog to the current EarthRadar tests, an Ohio State University radar has been tested under the ESTCP as a cued target-identification system. The radar would interrogate discrete points identified by magnetometers or EMI sensors as containing potential targets. The demonstration has shown the detectability of many targets in favorable conditions (homogeneous sand) on a gridded site at Tyndall AFB. The radar visited grid squares containing emplaced ordnance, emplaced clutter, or nothing. When discrimination of clutter from ordnance was attempted, the receiver operating characteristic (ROC) curve showed some capability above chance. For example, the best results were that a P_d of 0.9 was achieved with 0.5 P_{fa} (Ref. 20).

4.2   EarthRadar DATA DESCRIPTION

Reference 1 provides the contractor's results and supporting data for the EarthRadar testing at Fort A.P. Hill. That report, which is a large PowerPoint document, gives the assessed results for the target squares that were analyzed. Table 1 summarizes those reported results.

Table 1. Reported EarthRadar Results

Category	Number	Comments
Total blind squares	260	13 lanes by 20 squares/lane
Squares reported	112	2 passes not collected on all squares
Results reported:		
UXO	18	Assessed ID and depth provided in all squares
Clutter	65	Assessed depth provided in 20 squares
Empty	29	

Note that declarations were made on about 43 percent of the total available blind squares. There are two reasons for the incomplete results. First, the data-collection phase went much longer than planned or budgeted because of a combination of problems with the EMC that extended data collection and also because each square had to be traversed twice at relatively slow speed. Second, the analysis that leads to a declaration of square contents is involved and labor-intensive. For that reason, funding would not have allowed analysis of all the squares, even if data collection had been completed within the allotted time.

For each of the squares declared to contain a UXO target, a series of plots is provided in Reference 1. These include plots similar to Figures 5, 6, 7, and 8. The 3–D plots are labeled with the assessed UXO type. Squares assessed to contain clutter may or may not have plots provided. Only 20 of the 65 cells declared to contain clutter have a depth for the clutter indicated. Thus, the analysis of results is based on a relatively small sample. The actual distribution of UXO, clutter, and empty (no emplaced UXO or clutter) in the squares reported is representative of the entire site, however, and so while the sample size is small, the conclusions drawn concerning radar performance should be valid, within the limits of the statistics.

4.3 EarthRadar PERFORMANCE ANALYSIS

This subsection discusses the performance achieved by the EarthRadar in the blind testing. To protect the status of the pilot site for further blind testing, the results are generally couched in terms of percentages in this report. Table 2 summarizes the principal statistics of interest.

As noted in Section 4.1, the first requirement for UXO sensors is that they reliably respond to all potential UXO objects, whether or not they are able to discriminate between UXO and clutter. For these tests, that would be given by the probability that a

Table 2. EarthRadar Performance Statistics

Parameter	Definition	Value (%)
Probability of response	(number of UXO or clutter squares declared either UXO or clutter)/ (total number of UXO plus clutter squares)	73
Probability of false response	(number of empty squares declared UXO or clutter)/ (total number of empty squares)	75
P_d (UXO declared UXO)	(number of UXO squares declared UXO)/ (total number of UXO squares)	33
Probability that a UXO square was declared empty	(number of UXO squares declared empty)/ (total number of UXO squares)	29
Probability that a UXO square was declared clutter	(number of UXO squares declared clutter)/ (total number of UXO squares)	38
P_{fa} (empty or clutter squares declared UXO)	(number of empty plus clutter squares declared UXO)/ (total number of empty plus clutter squares)	14

square containing an emplaced object (UXO or clutter) is declared to contain either. That probability for the EarthRadar in these tests is 73 percent, not too far below the 85-percent requirement for the Kaho'olawe Tier II cleanup, but far below the near-100-percent value desired by regulators. A 100-percent response rate can be achieved, however, simply by declaring every square to contain either UXO or clutter. Therefore, in analyzing the probability of response, we must also consider how the declarations in UXO/clutter squares differed from those for empty squares. In this case, the numbers are essentially identical. That is, the EarthRadar declares the same fraction of empty squares to contain either UXO or clutter as it does the fraction of squares that actually do contain an object. If this data point were plotted on a ROC curve for response, it would fall on the chance diagonal. That is, the sensor shows no more capability to determine whether a square contains an emplaced object than a random guess provides.

Because of the time required to collect and process data, the EarthRadar does not appear to be a good candidate for a wide-area search sensor. Thus, it might more often be used to interrogate areas where another sensor has provided an alarm to determine whether the alarm is due to UXO or to clutter. In that mode, P_d, the ability to separate UXO from clutter, becomes important. The P_d achieved by the EarthRadar in these tests is 33 percent, far below values of interest to any of the communities involved in UXO clearance. The radar does show modest capabilities, as the P_d, combined with the P_{fa}, falls above the chance diagonal, indicating that self-clutter effects are occasionally overcome. Nevertheless, if we look only at squares containing UXO, the declarations appear random, with approximately a one-third probability that a square will be declared

to contain UXO, clutter, or be empty. Although the measured results differ somewhat from the nominal one-third value, recognize that they represent a single sample set from a random process. If the statistics for the underlying random process are calculated, the observed values are within 0.3 standard deviations of the most probable values of one-third each, given the number of samples.

Because accurate range or depth determination is one of the strengths of radar, one additional way to determine whether the EarthRadar is generally making declarations based on returns from emplaced objects rather than on system self-clutter is to look at the depth of the declared objects relative to their actual burial depth. A GPR that is detecting targets should produce a near-Gaussian-shaped depth-error distribution with a mean that is close to zero error. Such a distribution arises from an assumption of a number of uncorrelated, zero-mean, random error sources and application of the central limit theorem (Ref. 21). The difference of the mean from zero is a measure of a systematic depth offset (e.g., an error caused by use of an incorrect propagation velocity or a timing offset), and the standard deviation of the distribution is a measure of radar accuracy. If the depth errors (defined as actual depth minus declared depth) produce a uniform distribution, it is typically an indication that the radar is not detecting the actual targets. That is, if the target depths are uniformly distributed and the depth guesses are also uniformly random—not related to the actual target depth—a uniform depth error distribution results.

Figure 9 provides a depth-error histogram for all the EarthRadar blind test squares where a UXO or clutter object was declared, with an associated depth, and where the square actually contained an emplaced object, either UXO or clutter. There were only 20 instances where that occurred, so the histogram is sparse. Note, however, that the shape is close to a uniform distribution, tending to indicate no correlation between the radar's declared depth and the actual target depth.

Depth errors are signed, so average error is not a good measure of system accuracy because large negative and positive errors can offset each other to give a small mean error. More typically, depth-error standard deviation is used as the measure of GPR depth-accuracy capabilities, where the sample standard deviation is defined as (Ref. 22)

$$s = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(x_i - \bar{x}\right)^2} \quad , \tag{5}$$

Figure 9. Histogram of Depth Errors

with n the number of samples, x_i the depth error for the ith sample, and $\bar{x}$ the sample mean. Based on Equation 5, the depth-error standard deviation for the EarthRadar data in Figure 9 is 10.2 inches.

In assessing system performance, it is useful to compare the measured depth error standard deviation with one that would be obtained through random guesses at target depths. To make that comparison, IDA generated a Monte Carlo simulation using the actual depths of the clutter and UXO in the squares constituting the data in Figure 9. The Monte Carlo simulation generates synthetic depth error data using the following procedure:

- For each of the 20 cells, generate a uniformly distributed, random depth guess between 3 and 27 in. (The bounds of 3 and 27 in. were chosen because these were the minimum and maximum depth declarations made by the EarthRadar system in the 20 squares used to create Figure 9.)

- Compute the depth errors for the random guesses in the 20 cells.

- Compute the depth error standard deviation for the 20 samples.

- Repeat the procedure 500 times.

Figure 10 plots the depth-error standard deviation histogram for the 500 Monte Carlo trials. The histogram has a mean value of 11.1 in. and a standard deviation of 1.4 in. If we note where the EarthRadar's depth-error standard deviation of 10.2 in. would plot on the histogram, we find it is 0.9 in. smaller than the average obtained from the Monte Carlo, well within one standard deviation. Hence, the EarthRadar depth error

Figure 10. Histogram of Depth Error Standard Deviations for Repeated Trials Matching Reported EarthRadar Results

performance is consistent with the results for a system making random guesses, or at best, a system making mostly random guesses, with only an occasional depth call associated with an actual target.

5. SIGNAL PROCESSING AND EMC ANALYSES

5.1 SIGNAL–PROCESSING EVALUATION

The government asked IDA to assess the signal processing used by the EarthRadar. Preceding sections have provided an explanation of that processing and have discussed the results achieved. The purpose of this section is to evaluate the processing in an effort to assist understanding the system's capabilities and limitations.

The EarthRadar signal processing employs standard digital signal-processing techniques. In fact, most of the signal-processing functions used are implemented in the commercial packages that form the basis of the radar software, rather than in special software written by the radar developers. For example, the IFFT and window functions are provided in the Lab Windows CVI software. The 3–D model creation and manipulation are provided in the RockWorks 99 software. The one unusual aspect of the signal processing is the choice of the real part of the IFFT, rather than the magnitude, for target detection and imaging. As noted earlier, that choice is felt to be suboptimum because of the effect of the carrier signal modulation on the shape of the target response.

As noted in the Executive Summary and in the Data Description and Results sections, the response and detection performance of the EarthRadar were poor in these blind tests. That performance is not primarily a result of the signal processing, but of some well-known problems in GPR design and use that either have not been recognized or have been dealt with unsuccessfully. Each of those problems or limitations is described below.

5.1.1 Antenna Direct Coupling

Note the dark red and black bands beginning at the top of the display in Figure 6. Those bands represent a strong coupling directly between the transmit and receive antennas. Each half cycle spans about 20 cm in depth, which makes detection of small UXO objects at those depths very difficult. The direct antenna coupling signal appears in all the EarthRadar data. It is a well-known problem in GPR design. Two different approaches have historically been used to combat it. The first is to employ orthogonal polarizations for the transmit and receive antennas. This reduces the antenna coupling,

and as long as the target significantly depolarizes the incident wave, it increases the signal-to-interference ratio. A second approach, sometimes combined with the first, is to estimate and subtract the direct-coupling signal. This has been implemented by other researchers with some success, but changes in the underlying terrain also cause changes in the direct coupling, which limits the amount of cancellation that can practically be achieved. The EarthRadar uses neither approach and so is highly susceptible to performance degradation due to direct coupling.

5.1.2   Antenna Multiple Reflections

The bright and dark bands on Figures 6 and 7 indicate that electromagnetic energy is repeatedly reflecting between the two antennas, producing system self-clutter. At each reflection, a portion of the energy is absorbed and a portion is re-reflected. Because a radar measures distance by measuring transit time between transmission and reception, subsequent reflections appear progressively deeper. As the network absorbs some portion of the energy on each reflection, the signal also becomes progressively weaker. This coupling serves not only to potentially hide targets of interest, but also to help produce apparent targets. Although a particular artifact that appears in Figure 6 is circled as the selected target, note that similar features appear a number of locations on the color plot, as do other features that might just as easily be declared targets.

Figure 11, taken from calibration lane testing, illustrates both the problem of direct coupling and that of multiple reflections. IDA processed the stepped-frequency return from a single representative trace near the middle of an empty square. Because it contains no discrete clutter, the return from an empty square would be expected to smoothly decrease at a rate determined by the $1/R^4$ range falloff combined with soil attenuation. For this example, the relative amplitude of the radar return is plotted as a function of depth (IFFT bin number). In this case, each bin represents about a 2.8-cm depth difference in the earth (two-way return) or about 10 cm (one-way) in the cables connecting the network analyzer to the antennas. The first 100 bins represent the signal transit time through the two 5-m long cables between the network analyzer and antennas. EarthRadar processing removes this section of the return automatically, but we have chosen to display it in our processing. The very strong peak centered on bin 100 is the direct antenna coupling.

Bin 100 effectively locates the position of the surface in this plot. If we focus on the region between bins 100 and 200 (surface and a depth of about 2.8 m), we see a

**Figure 11. Magnitude of the IFFT of the Return from a Single Trace
from an Empty Calibration Square**

number of small, closely spaced peaks that represent a combination of reflections between the system antennas and IFFT processing sidelobes related to those reflections. The large peak near bin 200 is second-time-around cable return. That is, it is a coherent sum of the initial mismatch reflections at the transmit antenna and the receive port that have traversed the cables a second time. This artifact is visible in all of the EarthRadar data processed by IDA. Note that the noise floor of the network analyzer appears to be somewhat below −60 dB on this scale. System reflection returns in the depth region of interest for these tests are always at least 30 dB above the noise floor, and often more than 40 dB. Such strong components make it difficult to detect targets of interest, particularly if their return is small relative to the system self-clutter.

The problems provided by coupling and repeated reflections between the antennas are further illustrated in Figure 12, which is a color plot from a square, declared in Reference 1 to contain UXO, but which is actually empty. Note that changes in coupling between the antennas as the radar moves and the underlying ground surface varies change the character of the self-clutter. The result is to provide a large number of regions
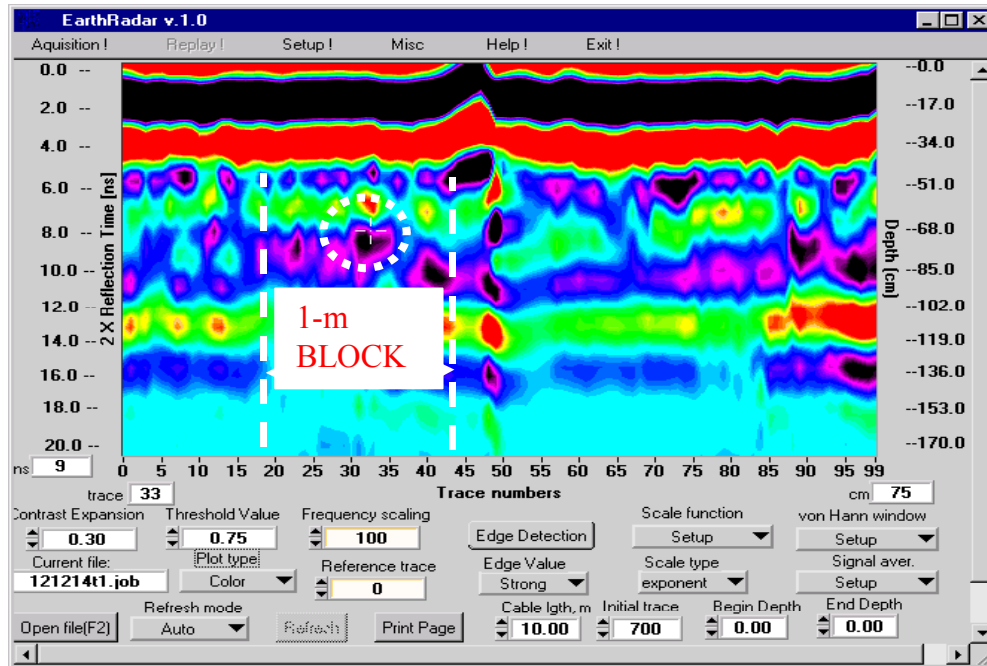
Figure 12. Color Plot from an Empty Square Declared to Contain UXO (Ref. 1)

containing target-like returns. This plot is typical and demonstrates the problem an operator would have in selecting an actual target return from among the returns due to system coupling.

5.1.3 Fundamental GPR Limitations

The process used by the EarthRadar operator to declare detections is to evaluate an image on a 3–D plot and assess its resemblance to a target of interest. Both shape and dimensions are used in the evaluation. There are several limitations imposed by the physics of the problem and by the processing used that make image analysis an unlikely avenue for successful target detection. These include resolution limitations in range and cross range and the scattering characteristics of targets in the frequency regime suitable for GPR operation. Each is discussed in turn.

5.1.3.1 Range and Cross-Range Resolution

The range resolution limitation for any radar is determined by the absolute bandwidth employed by the radar. The applicable equations are discussed in Section 3.3. For the 200-MHz bandwidth employed in the Fort A.P. Hill testing, the fundamental limit on resolving separate scatterers in range is that they be separated by just over 20 cm, so dimensions in the depth direction smaller than that would be nearly impossible to

5-4

ascertain. The EarthRadar does not use any synthetic-aperture processing. Therefore, the resolutions available in the two cross-range directions are dependent on the overlaps of the antenna beams in each direction and the depth of the target. Although patterns have not been measured for the antennas used, dipole patterns are typically very broad, so it seems unlikely that cross-range resolutions would be capable of supporting finely gridded 3–D images.

5.1.3.2 Electromagnetic Scattering Characteristics

A more fundamental physics limitation argues against the use of 3–D images for UXO detection and discrimination. That is, at the wavelengths useful for earth penetration, the targets of interest do not scatter electromagnetic energy diffusely. Our eyes use the diffuse scatter from the entire surface of an object at optical frequencies to ascertain the object's shape. Surface roughness must be a significant fraction of a wavelength to support diffuse scatter (at normal incidence, a tenth wavelength RMS surface roughness causes about 80 percent of the incident power to be scattered diffusely). At optical frequencies the wavelength is between 0.4 and 0.7 μm, so a very small roughness is required for diffuse scatter from UXO objects. In this set of tests, the EarthRadar operated at frequencies below 1 GHz. Even with an index of refraction of 3.5, the wavelength in the ground was 8.5 cm or greater, so diffuse scatter from UXO objects was not supported. In general, scattering at GPR frequencies from UXO and clutter is dominated by specular returns. That is, scattering is received from portions of the body where the angle of incidence and angle of reflection are equal. As this is a near-monostatic geometry, the specular requirement translates to dominant scattering from surfaces and edges whose normals point back toward the radar. For the typical target, only a small portion of the surface will support specular scatter, and so it will be very difficult to map the return into the shape providing it.

5.2 EMC ANALYSES

On several occasions during the data collection, a disabling interference was observed between two components of the EarthRadar system. Specifically, the network analyzer generated harmonics and spurious frequencies within the RF bandwidth of the on-sled GPS antenna and its receiver at a level sufficient to prevent GPS function. The following paragraphs describe the rationale for this concern and outline the observations made during EarthRadar data collections at Fort A.P. Hill.

5.2.1 Analysis

Reception of the coarse acquisition (C/A) code modulated onto the L1 carrier at 1,575.42 MHz is essential for successful GPS operation. This signal level varies with satellite elevation and range, but it can be as low as –130 dBm measured at the output of an antenna with 0-dB gain referenced to an ideal, isotropic, circularly polarized antenna. On the other hand, the maximum level of harmonic and spurious emissions allowed by the performance specification of the HP 8753D network analyzer used within the EarthRadar is given in Table 3, where dBc indicated decibels referenced to the carrier power. The values in Table 3 were taken from the HP 8753E User's Guide (Ref. 23), but Agilent Technologies technical support personnel confirmed that specifications are identical for the D model. Note that the maximum power output of the network analyzer is 10 dBm and can be adjusted over the range –85 dBm to 10 dBm. Dr. Bakhtar has indicated that normal practice is to operate the analyzer at maximum power, which would produce the highest possible signal-to-receiver-noise ratio within the EarthRadar. This would normally be desirable, if the system performance of the EarthRadar were noise limited.

Table 3. Signal Purity from 8753D Network Analyzer (Ref. 23)

	2nd Harmonic	3rd Harmonic	Nonharmonic Spurious
at Maximum power	–25 dBc	–25 dBc	–30 dBc
10 dB below max power	–40 dBc	–40 dBc	NA
20 dB below max power	–50 dBc	–50 dBc	–55 dBc

Note that the harmonic and spurious output decreases dramatically with reduced output power. A 10-dB reduction in system output power would reduce absolute harmonic power by 25 dB. Table 4 provides a rough, order-of-magnitude estimate of the potential power at the GPS receiver input due to the harmonic/spurious output of the network analyzer. This table assumes that the network analyzer is producing harmonics at the maximum level allowed by its specifications in Table 3.

The system-specified minimum C/A code signal level is –130 dBm, which is further enhanced by 30 dB of C/A code processing gain. This leads to a net potential GPS signal-to-interference ratio of –11 dB, which would be unacceptable for proper C/A code system operation. With the exception of the harmonic power output of the network analyzer and the propagation loss, we note again that all of the parameters in Table 4 are order-of-magnitude estimates. Also note that –130 dBm is a minimum signal power level

Table 4. Estimated Available Interference Power at GPS Receiver

	Gain (+)	Loss(–)
Harmonic/spurious power available at EarthRadar transmission antenna input	–15 dBm	
Mismatch loss at EarthRadar antenna terminals		3 dB
EarthRadar antenna gain in the direction of the GPS	–10 dBi	
Propagation loss from radar to GPS antenna		21 dB
Antenna choke loss		20 dB
GPS antenna backlobe gain	–20 dBi	
Estimated interference power available at GPS receiver	–89 dBm	

that is often exceeded in practice, particularly when satellites are well above the horizon. The actual interference power within the GPS receiver passband will depend on the particular choice of EarthRadar frequencies as well as the receiver selectivity characteristics. Although crude, this analysis shows the potential for a serious problem and is consistent with much of the anecdotal evidence associated with EarthRadar operability problems.

The HP8753D network analyzer generates a continuous-frequency sweep that begins somewhat below the requested "start" frequency (to allow the sweep rate and signal level to stabilize) and increases to the requested "stop" frequency. Anytime the swept band includes a subharmonic of the L1 or L2 GPS frequency there is a potential for serious self-interference (see Table 5).

Table 5. GPS Frequencies and Subharmonics

Satellite signal	Fundamental/MHz	2nd Subharmonic/MHz	3rd Subharmonic/MHz
L1	1,575.42	787.71	525.14
L2	1,227.6	613.8	409.2

5.2.2 Testing and Results

Since the specific network analyzer used in the EarthRadar was heavily committed to the data-collection program, it was deemed impractical to remove it to a laboratory for RF measurements. Under the circumstances, the most practical way to validate the specific interference hypothesis was by direct observation of the phenomenon in the field. That is, we observed how GPS functioning changed with the frequency range of the EarthRadar. The following observations were made at Fort A.P. Hill on 14 March 2001:

- When the EarthRadar sweep included a second or third subharmonic of L1 (see Table 5), the GPS receiver was consistently unable to maintain satellite lock. This effectively disabled the data-collection process.

- When the EarthRadar sweep was adjusted to begin above 800 MHz (excluding 2nd or 3rd subharmonics of L1) the GPS receiver functioned normally.

These observations were repeated consistently over a period of ~1 hour using several orientations of the EarthRadar and GPS antennas. In addition, the system was tested with a low-pass filter inserted at the output of the network analyzer. The characteristics of this filter are shown in Figure 13. This filter provides ~50 dB of suppression at the L1 frequency. With this filter inserted between the network analyzer and the transmission antenna, the EarthRadar was set to sweep from 500 to 1,000 MHz. No loss of GPS performance was observed. Based on these field observations, which are consistent with the order-of-magnitude analysis performed above, we conclude that the operation problems encountered during the course of the EarthRadar data collection were due to self-interference between harmonics of the transmitted signal and the GPS receiver used to locate the EarthRadar antenna.


Figure 13. Insertion and Return Loss

6. CONCLUSIONS

6.1 GENERAL

Based on the results of the Fort A.P. Hill testing reported in Reference 1 and analyzed in Section 4, we reach the following conclusions:

- *The EarthRadar showed no capability to provide a preferential response to squares containing emplaced UXO or clutter over empty squares.*

  This conclusion is supported by the nearly identical response rate for squares with emplaced objects (73 percent) and empty squares (75 percent).

- *In squares containing UXO, the EarthRadar declarations were consistent with random guesses among the three possibilities.*

  This conclusion is supported by the detection performance, where the declarations in squares that actually contained UXO appeared to be equally distributed (33 percent correctly declared UXO, 38 percent incorrectly declared clutter, and 29 percent incorrectly declared empty).

- *Most, if not all, of the UXO and clutter declarations were based on self-clutter signals generated by the EarthRadar, rather than on returns from actual targets.*

  This conclusion is supported by the response and detection performance discussed above, but is bolstered by the depth-error statistics discussed in Section 4.3. The fact that the depth-error standard deviation for EarthRadar declarations closely resembles that provided by random depth guesses strongly implies that the radar detections have no relation to the actual targets.

- *The EarthRadar concept for target detection and identification is fundamentally flawed.*

  The EarthRadar depends on diffuse scatter from targets and for sub-wavelength resolution for image formation. In the frequency bands available to GPR, UXO targets scatter specularly. Neither electromagnetic theory nor signal-processing laws support obtaining the required range and cross-range resolutions, given the frequencies and bandwidths employed by the EarthRadar. Arbitrary and subjective scaling and thresholding allow the

operator to create object-like images from the data, but there is no evidence that image shape and original data have any strong correlation. That most of the declarations in these tests appeared to be based on system-generated self-clutter and the poor performance in the tests support the hypothesis that the 3–D images produced have little or no relationship to scattering objects beneath the radar antennas.

6.2 REASONS FOR POOR PERFORMANCE

As noted in Subsection 5.1, the reasons for the poor performance stem partly from system design problems and partly from erroneous assumptions concerning target-scattering characteristics at GPR frequencies. In summary, they are as follows:

- Strong direct coupling between the transmit and receive antennas hides shallowly buried targets;

- Multiple reflections between the antennas create system self-clutter that can either hide targets or be mistaken for targets;

- The bandwidth employed and the real-aperture processing do not achieve resolutions in any dimension sufficient to provide accurate size information on many of the targets; and

- The imaging techniques used for target detection/identification depend on target diffuse scatter so that an image of the target's geometrical shape will be formed; that is almost never the case for the targets of interest at GPR frequencies. It is definitely not the case for the targets and wavelengths employed in this testing.

6.3 OTHER SHORTCOMINGS

Even if the system had shown good UXO detection capabilities, there are a number of areas in which the system would have to show marked improvements before it would be a useful tool in the field. These include the following:

- Data collection is extremely slow, less than 1 m^2/min under best conditions. In the Fort A.P. Hill testing, effective collection rates were much slower than that.

- Data processing is also slow and is labor intensive. Even for the case where target location was known exactly, data manipulation, plot and 3–D image production, and target detection/classification required several hours of labor per target.

- Target selection is highly subjective. Plots shown to the government team always appeared to contain multiple possible targets. The EarthRadar contractor was unable to provide objective criteria for the selection of a

particular plot anomaly as a target, rather than some other, very similar anomaly, except for its location near the center of a grid square. Based on the test results, it appears that the EarthRadar contractor is unable to successfully separate system artifacts from targets of interest.

6.4 SUMMATION

All of the problems affecting EarthRadar performance—direct antenna coupling, self-clutter, limits on resolution, and target-scattering characteristics—are well known to the GPR community. Solutions have been posited for some of them (antenna coupling and self-clutter) and implemented by other researchers with varying degrees of success. The limitations on resolution and the scattering characteristics of targets at the frequencies that can be employed by a GPR, however, provide fundamental limitations on the signals available for processing. GPRs must work within those limitations to detect and identify buried objects. Other investigators are exploring synthetic-aperture processing, target polarimetric and resonance characteristics, and sophisticated target recognition algorithms to improve performance. Nevertheless, radar remains an unproven sensor for the UXO detection problem. The EarthRadar, based on a fundamentally flawed hypothesis concerning target scattering characteristics and image-formation capabilities, does not offer any technology improvements that change the current status.

GLOSSARY

3–D	three dimensional
AC	alternating current
ARC	active range clearance
C/A	coarse acquisition
DC	direct current
EMC	electromagnetic compatibility
EMI	electromagnetic induction
EOD	Explosive Ordnance Disposal
ESTCP	Environmental Security Technology Certification Program
FM	frequency modulated
GPR	ground-penetration radar
GPS	Global Positioning System
GUI	graphical user interface
IDA	Institute for Defense Analyses
IDFT	inverse discrete Fourier transform
IFFT	inverse fast Fourier transform
JUXOCO	Joint Unexploded Ordnance Coordination Office
MTADS	Multisensor Towed Array Detection System
R&D	research and development
RFI	radio frequency interference
ROC	receiver operating characteristic
RTI	range-time-intensity
SBIR	Small Business Innovative Research
UHF	ultra high frequency
UXO	unexploded ordnance
VHF	very high frequency

REFERENCES

1. Bakhtar, K., "EarthRadar Evaluation Tests at Fort A.P. Hill, Virginia," Final Report on Contract F08630-98-C-0031, Bakhtar Associates, 3420 Via Oporto, Suite 201, Newport Beach, CA 92663, 23 May 2001.

2. Adams, J.B., "Unexploded Ordnance Recovery Depth Database," U.S. Army Engineering and Support Center, Huntsville, Ala.

3. "Understanding the Fundamental Principles of Vector Network Analysis," Application Note AN 1287-1, Agilent Technologies.

4. Lane, T.L., "Stepped-Frequency and ISAR Imaging Systems," in *Coherent Radar Performance Estimation,* J.A. Scheer and J.L. Kurtz, editors, Artech House, Boston, 1993, pp. 234–244.

5. Bakhtar, K., E. Sagal, and J. Jenus, "Detection and Discrimination of Buried UXO at the Naval Weapons Station Using U.S. Air Force EarthRadar," 29th Department of Defense Explosive Safety Seminar, New Orleans, La., 18–20 July 2000.

6. Curtis, J., "Dielectric Properties of Soils, JUXOCO Pilot Site, Fort A.P. Hill," Data Report, U.S. Army Corps of Engineers, 8 September 1998.

7. Harris, F.J., "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform," *Proceedings of the IEEE*, Vol. 66, No. 1, January 1978, pp. 51–83.

8. Porat, B., *A Course in Digital Signal Processing,* John Wiley & Sons, New York, 1997, pp. 164–168.

9. "Report of the Defense Science Board Task Force on Unexploded Ordnance Clearance, Active Range Clearance, and Explosive Ordnance Disposal Programs," April 1998.

10. Kaho'olawe Clearance Project Brief, Kaho'olawe Clearance Conference, November 1999.

11. Environmental Requirement A (1.6a), Unexploded Ordnance Screening, Detection and Discrimination, Draft.

12. "Handheld Gradiometer Survey Test at the Marine Corps Air Ground Combat Center," Twentynine Palms, Calif., NAVEODTECHCENTER, September 1992.

13. "Advanced UXO Detection Discrimination Technology Demonstration—U.S. Army Jefferson Proving Ground," Madison, Ind., 2001.

14. "Multi-sensor Towed Array Detection System," ESTCP Cost and Performance Report, September 1999.

15. McDonald, J., H.H. Nelson, B. Puc, "MTADS Magnetometer and EM Surveys at Fort Ord," Naval Research Laboratory Publication NRL/PU/6110-010349, March 5, 2001.

16. Nelson, H.H., and B. Barrow, "Electromagnetic Induction and Magnetic Sensor Fusion for Enhanced UXO Target Classification," Naval Research Laboratory Publication NRL/PU/6110-00-423, 29 June 2000.

17. Altshuler, T.W., A.M. Andrews, R. Dugan, V. George, M. Mulqueen, D. Sparrow, *Demonstrator Performance at the Unexploded Ordnance Advanced Technology Demonstration at Jefferson Proving Ground (Phase I), and Implications for UXO Clearance,* IDA Paper P-3114, October 1995.

18. "Unexploded Ordnance Advanced Technology Demonstration at Jefferson Proving Ground (Phase II)," Report SFIM-AEC-ET-CR-96170, June 1996.

19. Sichina, J., M. Ressler, A. Sullivan, "Low Frequency Ultra Wideband Synthetic Aperture Radar for Remote Detection of UXO," SERDP In Progress Review, April 2000.

20. Chen, C.C., M.B. Higgins, K. O'Neill, "Data Processing Results for UXO Classification Using UWB Full-Polarization GPR System, Tyndall AFB Site Demo," 2000.

21. Papoulis, A., *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill, New York, 1965, p. 266.

22. Vardeman, S. B., *Statistics for Engineering Problem Solving*, IEEE Press, New York, 1994, p. 81.

23. "User's Guide: Agilent Technologies 8753ET and 8753ES Network Analyzers," Part Number 08753-90472, Agilent Technologies, May 2000.

REPORT DOCUMENTATION PAGE

Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE September 2001	3. REPORT TYPE AND DATES COVERED Final — June 2000 – July 2001

4. TITLE AND SUBTITLE
Evaluation of EarthRadar UXO Testing at Fort A.P. Hill

5. FUNDING NUMBERS
DASW01 98 C 0067
AK-2-1923

6. AUTHOR(S)
M.T. Tuley, J.M. Ralston, F.S. Rotondo, A.M. Andrews, E.M. Rosen

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311-1772

8. PERFORMING ORGANIZATION REPORT NUMBER
IDA Document D-2625

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
ODUSD(S&T)
Suite 9030
1777 N. Kent Street
Rosslyn, VA 22209

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 180 words)*

This report provides an analysis of EarthRadar blind testing carried out at the Joint Unexploded Ordnance Coordination Office (JUXOCO) Pilot Site, Fort A.P. Hill, Va., during the fall of 2000 and spring of 2001. The EarthRadar, developed by Bakhtar Associates (Newport Beach, Calif.) under an Air Force Small Business Innovative Research contract, is intended to detect and classify buried objects. The report describes the JUXOCO test facility; the EarthRadar hardware, signal processing, and data analysis; and the blind lane test results. It evaluates EarthRadar performance in the context of documented requirements for systems employed in UXO detection/discrimination and draws conclusions regarding the limitations of this technology for UXO clearance applications.

14. SUBJECT TERMS
ground-penetrating radar, UXO, target ID

15. NUMBER OF PAGES
57

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR